GRANT BOLLMER

Perplexing Intelligence

AI and the Aesthetics of Statistics

ABSTRACT This essay argues that the experience of perplexity, as discussed in this special issue, is negated or foreclosed by the statistical measures used in artificial intelligence—which includes, interestingly enough, a specific measure of uncertainty termed *perplexity.* This argument proceeds in two parts. The first reviews recent discussions of generative AI to argue that current versions of AI are distinct from the symbolic logic that guided earlier forms of AI and instead create an *environment* for human experience that cannot itself be grasped by human sensation. The inequivalence between human sensation and the technical foundations of this "world" results in the many hyperbolic and hysteric discussions of the potentials and problems of artificial intelligence today. The second part details the technical specifics of how statistical logic is implemented in well-known generative AI systems. It follows how generative AI derives from the measures of probability first described in the information theory of Claude Shannon and discusses specific datasets used to train and manage the "perplexity" of specific AI models, including ChatGPT and DALL-E. This essay concludes with a brief discussion of how the technical operations of generative AI relate to a distinction between words, images, and numbers, and how this distinction—or conflation—demonstrates that, in some way, human beings are aware that a world or environment shaped and conditioned by generative AI is, effectively, at odds with and exists beyond human comprehension.
KEYWORDS artificial intelligence, statistics, aesthetics, thought, symbolic logic, medium specificity, perplexity

In thinking through the claims of this special issue's introduction,[1] one way we can define an aesthetics of perplexity is as a sensory mismatch or inequivalence. The qualities of an object perceived by an individual are distinct from the object itself. The experience of the object, rather than enabling a transmission between perceiver and perceived, opens a gap that appears, at first glance, to manifest in bafflement, uncertainty, incomprehension. This perplexity is a confrontation, though it isn't a hostile confrontation and never ends with the viewer's perception "corrected" through the encounter with the object.[2] Understanding is never really achieved, means never resolve into ends. Instead, the experience of perplexity is a suspended state of potentiality that refuses clear answers or any settling into accepted categories of sense-making.

1. Nina Peterson and Katherine Guinness, "Introduction," this issue.
2. See the discussion of the "orthopedic aesthetic" in Grant H. Kester, *Conversation Pieces: Community and Communication in Modern Art* (University of California Press, 2004).

Afterimage, Vol. 52, Number 1, pps. 111–129. ISSN 2578-8531. DOI: https://doi.org/10.1525/aft.2025.52.1.111

Much aesthetic experience of digital culture could be framed in these terms, from memes that are funny because of their nonsensical nature, through the varieties of unusual internet stories found on Reddit, to the extremely bizarre and disturbing images of "DeepTok." But "weird" is more often the term used to describe the stranger, more confounding aspects of digital media. As Mark Fisher has argued, the weird as an aesthetic category relates to "a fascination for the outside, for that which lies beyond standard perception, cognition, and experience. This fascination usually involves a certain apprehension, perhaps even dread—but it would be wrong to say that the weird [is] necessarily terrifying."[3] From the stories of H. P. Lovecraft to the music of The Fall, for Fisher the weird is about the presence of something "*which does not belong.*"[4]

The perplexing, I would suggest, is similar—it draws attention to something that cannot fully be assimilated to oneself and one's sensory or cognitive abilities, though not inherently in a way that "does not belong." Rather, the perplexing is undecidable, and may signal that the *viewer* is the one who does not belong, that the world may not be for me or you or us, and the confusion or bafflement that appears as perplexity points to this absence of alignment.[5] When it comes to digital culture, the weird may be vanishing—even though perplexity is not—because of the overwhelming preponderance of images and text generated by artificial intelligence, the rise of "AI slop" increasingly saturating online space. With digital media, perplexity is a category that refers to a feeling that digital culture and the internet do not inherently require the presence of human beings. What I suggest in this essay is that the ability to sense and apprehend an outside, when it comes to the internet and digital culture—which manifests itself in one sense of what we mean by perplexity—is increasingly foreclosed beyond a sneaking suspicion that the proper location of humanity is in this outside. In other words, AI is perplexing because we humans are the weird ones on the internet, not the bots.[6]

To make this claim, I draw on another, second definition of perplexity, a definition that is commonly used in research on AI. At first glance, the images and text generated by AI would seem to be exceptionally "weird." As Chris Chesher and César Albarrán-Torres describe it, AI imagery is "uncanny" because "it presents vivid images of people who were never alive and places that never existed . . . [it] foster[s] in its users an assumption that all images are suspect, even if they seem to resemble someone and express emotions."[7] Obviously, there is some agency making decisions beyond human awareness, an outside

3. Mark Fisher, *The Weird and the Eerie* (Repeater, 2016), 8–9.

4. Fisher, *The Weird and the Eerie*, 10.

5. In this way, I also follow the introduction in how perplexity and the absurd share a number of presumptions.

6. It may appear that I'm creating a hard distinction between the human and the technological, though I believe that any definition of "the human" is co-constituted with "the technological." These are mutually emergent, ontogenic categories rather than stable "things." There are power differences in this co-constitution, however. My claims (and ontological assumptions) here are similar to those made by Bernard Stiegler in *Technics and Time, 2: Disorientation*, translated by Stephen Barker (Stanford University Press, 2008), which argues that this mutual emergence shifted in the twentieth century to devalue human experience relative to technical rationality. See also the arguments I make in *Materialist Media Theory: An Introduction* (Bloomsbury Academic, 2019).

7. Chris Chesher and César Albarrán-Torres, "The Emergence of Autolography: The 'Magical' Invocation of Images from Text through AI," *Media International Australia* 189, no. 1 (2023): 67.

that cannot and will never be fully brought to heel by human experience. In its very existence, AI text and images confront humanity with other agencies that are inherently beyond the human. But my sense is that these aren't inherently weird, or genuinely about some kind of outside otherness, because of the *statistical reason* on which generative AI relies—a statistical reason that has been shaping human existence, if with a wide range of variants and permutations, since at least World War II.

One statistical measure used to judge the quality of an AI system is literally termed "perplexity." What I describe in this essay is how a focus on the perceptual contents of AI systems made sensible to human beings—the text, the images—distract us from the real problem of artificial intelligence. Debates about whether or not AI can think or be creative, or if AI reasonably simulates human thought and perception, overlook the statistical basis of all AI.[8] In other words, what I argue here is that the aesthetics of perplexity, when it comes to AI, is less about what appears to human experience as words or text, and more about that which cannot and does not enter into the sensation of AI at work: the statistical logic of probability as built into the technical function of AI. Unlike the weird or the obviously perplexing, this statistical logic determines, but never directly enters into, human experience. What then does it mean to suggest an aesthetics of something that never can be directly sensed? In asking this question, I'm following the work of a scholar like Fabian Offert, who has argued that AI systems produce their own models of representation, and "we cannot trust [them] to represent the world in a way that is coherent to us."[9] Yet, while the relationship between (textual) input and (visual or textual) output observed by human users and researchers may have little to do with the actual operations of AI systems, this doesn't mean that the processes and epistemological presumptions that guide their functioning are completely "black boxed."[10] Trying to draw out some of these processes and presumptions is part of what this article attempts to do, even though I still embrace a view that suggests the impenetrability of AI (and thus the devaluing of human agency) is one objective of these systems.

My argument proceeds in two parts. The first half of this essay examines several analyses of generative AI and ultimately claims, following Luciana Parisi, that generative AI should be thought of as providing an *environment* for conditioning human sensation and experience that relies on statistical reason rather than logic or symbolic reference. The second part asks, If AI is shaping this sensory environment, on what basis does it do

8. Despite my focus here, all of these debates around AI are secondary to the catastrophic energy use inherent in its deployment. See Angely Mercado, "Microsoft Is Using a Hell of a Lot of Water to Flood the World With AI," *Gizmodo*, September 11, 2023, https://gizmodo.com/microsoft-water-usage-ai-iowa-data-center-1850826419; and Lauren Leffer, "The AI Boom Could Use a Shocking Amount of Electricity," *Scientific American*, October 13, 2023, www.scientificamerican.com/article/the-ai-boom-could-use-a-shocking-amount-of-electricity. On the general energy politics of digital technologies and data centers see also Mél Hogan, "Facebook Data Storage Centers as the Archive's Underbelly," *Television & New Media* 16, no. 1 (2015): 3–18, and Sy Taffel, "Data and Oil: Metaphor, Materiality and Metabolic Rifts," *New Media & Society* 25, no. 5 (2023): 980–98.

9. Fabian Offert, "Can We Read Neural Networks? Epistemic Implications of Two Historical Computer Science Papers," *American Literature* 95, no. 2 (2023): 425.

10. Fabian Offert and Ranjodh Singh Dhaliwal, "The Method of Critical AI Studies, A Propaedeutic," Cornell University Library, arXiv.org, doi.org/10.48550/arXiv.2411.18833, 4–6.

so? I go into detail about the statistical measure of perplexity, provide some examples of how perplexity works, and describe how it derives from the arguments of Shannon's "mathematical theory of communication." I conclude this essay with some speculations on how the statistical basis of AI relates to longstanding debates about aesthetics and medium specificity, and how the experience of AI not only conflates differences between images and text, but does so by reducing all human sensation to numerical "tokens," a level of reality that does not itself enter into human sensation. The ultimate point of this essay is to claim that if we are in some way aware of this agency determining the space within which humanity operates—and I think we are—then we are also in some way aware that this "world" of statistical reason, as produced through the mathematics of probability, is not really "for" human experience or agency. Humanity's proper location, with the perplexing experience of AI, is itself within the "outside" of digital culture. This is a conclusion with several contradictions, I should note, and perhaps a future direction for "critical AI" studies is to more fully draw out and embrace these contradictions rather than make definitive claims about what AI is and does—drawing out the perplexity inherent when confronted with nonhuman, statistical agencies in a world that seems to minimize human action.

GENERATIVE ARTIFICIAL INTELLIGENCE AS AN ENVIRONMENT FOR EXPERIENCE

I want to start by thinking through the specific political and aesthetic implications of AI today. The past decade has seen—quite understandably—an explosion in commentary on the social, technical, and political implications of generative artificial intelligence. The "generative" nature of today's AI differentiates it from earlier experiments with "symbolic" AI from the 1960s until the '90s, which are often dismissed with the folksy name "good old-fashioned AI" and its decidedly less folksy acronym GOFAI. These earlier AI models once suggested that computers could understand natural language and respond accordingly through a rationalist implementation of symbolic logic—a model of intelligence criticized by a range of philosophers, including Hubert Dreyfus and John Searle, using a range of arguments, some better than others.[11] The essential assumption here is that AI would "work" once a system could be developed to link words with objects, actions, and other words—to understand human language as based in symbolic reasoning and inference.

Debates around GOFAI set a path for countless philosophical and technical problems that persist today. A regular belief stemming from the intersection of AI and the philosophy of mind—named computationalism—positions an AI that passes as human (or passes what is colloquially named the "Turing Test," derived from mathematician Alan Turing's suggestion that a machine should be understood as thinking if it appears to be thinking) as providing a literal model for human cognition. With a successful AI, one could reverse engineer the functioning of human consciousness. Thus, computationalism

11. See Hubert L. Dreyfus, *What Computers Still Can't Do: A Critique of Artificial Reason* (MIT Press, 1992); John Searle, *Minds, Brains and Science* (Harvard University Press, 1984).

assumes that thinking, by humans and computers alike, is an essential capacity for logical reasoning through symbols, and that thinking is equivalent to the human sensation and projection of thinking onto technologies and other people alike.[12] The category of reference for determining "thought" is how humanity senses the performance of thinking, and the technical operations of a computer that appears to think are then retroactively projected onto the human mind. Today's AI often operates through a completely different means than symbolic AI—using neural networks that mimic synaptic functioning rather than logical reasoning, for instance—yet similar beliefs persist: if we can make an AI that appears as if it can think, then the technical operations of that AI can be used as a model for reverse engineering human cognition—a model that, importantly, is not metaphorical.

Much of the discussion surrounding generative AI, however, reveals a problem with this line of argumentation. Even though many of the most popular AI systems appear to "pass the Turing Test," almost all "successful" AI systems today represent the world in a way that has little to no relation with human sensation, language, and symbolic logic.[13] Generative AI turns away from a reliance on logic and employs algorithmic and probabilistic modeling, which is, as Luciana Parisi describes in an early critical work on generative AI, itself dependent on incomputable numbers such as infinity and what the mathematician Gregory Chiatin terms *Omega*.[14] Generative AI, for Parisi, is a major element of "postdigital" and "postcybernetic" culture, as

> generative algorithms are entering all logics of modeling—so much so that they now seem to be almost ubiquitous (from the modeling of urban infrastructures to the modeling of media networks, from the modeling of epidemics to the modeling of populations flows, work flows, and weather systems)—so too are their intrinsic incomputable quantities building immanent modes of thought and experience.[15]

The "generative" of generative AI is, as Parisi I think correctly identifies, not merely about the creation of something new or undetermined, but "the programming of events," where our current "postcybernetic culture is dominated not by the suprasensory bombardment of too much information, but by the algorithmic prehensions of incomputable data . . . "[16] These algorithmic prehensions are beyond human sensation, shaping and guiding human experience in space and time.[17] In this context, it is not so much that AI provides a model for the structure and function of human thinking, but instead provides an *environment* in which the possibilities and limitations of human thought are guided, an environment which

12. For the original description of the "Turing Test," see Alan M. Turing, "Computing Machinery and Intelligence," *Mind* 59, no. 236 (1950): 433–60. For an overview of computationalism and other similar philosophical positions see Michael Rescorla, "The Computational Theory of Mind," *Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta (Stanford University, 2024), plato.stanford.edu/entries/computational-mind.

13. Offert, "Can We Read Neural Networks?"

14. See, for example, Gregory Chiatin, *Meta Math! The Quest for Omega* (Knopf Doubleday, 2006).

15. Luciana Parisi, *Contagious Architecture: Computation, Aesthetics, and Space* (MIT Press, 2022), 18.

16. Parisi, *Contagious Architecture*, 79–80.

17. See Mark B.N. Hansen, *Feed-Forward: On the Future of Twenty-First-Century Media* (University of Chicago Press, 2014).

is itself beyond human sensation. When this broader environmental shaping of human sensation is implemented at an engineering level, as Offert and Thao Phan have shown, it manifests through attempts to correct "bias" by changing or augmenting the textual prompts of users, shifting from a "model-based to user-based debiasing" that assumes that AI models are "complete" or totalizing from the outset.[18] Or, the AI is presumed "correct" and the human users are the "problem" that leads to biased (racist, sexist) output. This is, I might add, a presumption with some irony, given how the potentials of generative AI also require a perpetual feed-forward toward the infinite and the incomplete.

Thus, generative AI emerges within a culture in which computation directs access to potentiality beyond empirical experience, potentiality shaped by the algorithmic and probabilistic interpretation of massive and theoretically infinite sets of data, in which decisions and actions are delegated to systems that remain partially or completely obscured to human beings. Given this relation to human knowledge and experience, it makes sense that much of the popular commentary about AI from the past several years tends to be utopian or dystopian, placing the survival or demise of humanity completely in the wake of what's often referred to as "artificial general intelligence," or AGI. While Parisi highlights how all forms of generative AI can be said to possess a kind of programmed agency beyond logical determinism, AGI identifies a form of AI that can act on its own in relation to scenarios that are completely new and unpredictable, a form of AI that, in the minds of many of those in the tech industry, has yet to emerge but is the ultimate goal of their labor.[19]

The rise in hyperbolic and often hysteric discussions of AI comes from an understandable problem: we are aware, in some way, of how AI is the environment for determining human sensation and cognition, and this environment for sensation and cognition seems limited and distorted given the perpetual "enshittification" of the contemporary internet.[20] More importantly, this environment is beyond human apprehension. "If we are worried about [AI] rendering humans obsolete," Leah Henrickson writes, "it is because somewhere within ourselves we already feel irrelevant."[21] Henrickson's claim resonates with one I made nearly a decade ago, in which I argued that social media—despite any marketing about "connecting people"—relies on a reductive and instrumental understanding of human capacity that ultimately renders human beings extraneous.[22]

But this feeling has been taken in an exceptionally broad number of directions, not all of which are genuinely "critical" of or even opposed to AI at all—for instance, what's often termed the "doomer" critique of AI uses the apocalyptic fantasy of uncontrolled AI

18. Fabian Offert and Thao Phan, "A Sign That Spells: DALL-E 2, Invisual Images and The Racial Politics of Feature Space," arXiv:2211.06323, doi.org/10.48550/arXiv.2211.06323.

19. See Louise Amoore, Alexander Campolo, Benjamin Jacobsen, and Ludovico Rella, "A world model: On the political logics of generative AI," *Political Geography* 113 (2024): 1, doi: www.sciencedirect.com/science/article/pii/S0962629824000830?via%3Dihub.

20. Cory Doctorow, "'Enshittification' is Coming for Absolutely Everything," *Financial Times*, February 8, 2024.

21. Leah Henrickson, "Conversations with No One," *Poetics Today* 45, no. 2 (2024): 297.

22. Grant Bollmer, *Inhuman Networks: Social Media and the Archaeology of Connection* (Bloomsbury Academic, 2016).

destroying the world to argue for the necessity of *perfecting* AI as quickly as possible. The manifold interpretations of what AI is or does rely on the simple fact that almost nobody knows, or at least seems to care about, what AI actually does in making its decisions, instead defaulting to science fictional narratives and surface-level effects, *even when it is possible to examine the general logics that underpin generative AI*—which is another point that differentiates GOFAI and its symbolic logic with generative AI. There is an exceptionally common presumption that it's not possible to understand what generative AI actually does, a presumption shared by countless humanists and technologists alike.[23] And this isn't even to mention that "artificial intelligence" itself refers to a massive number of distinct and irreducible technical operations, from machine learning[24] to machine vision.[25] For Emily Bender, the co-author of one of the first major critiques of generative AI, artificial intelligence is "a marketing term. It's a way to make certain kinds of automation sound sophisticated, powerful, or magical and as such it's a way to dodge accountability by making the machines sound like autonomous thinking entities rather than tools that are created and used by people and companies."[26]

Of all variants of AI, learning language models (LLMs) have been most scrutinized by those in the arts and humanities—perhaps because the capacity of ChatGPT and other text generators to produce "good" (or at least difficult to identify) texts best fulfills the Turing-inspired definition of thinking popular today, and thus provokes the most obvious questions about human creativity and thought. ChatGPT and its often impressive output is a seeming realization of a poststructuralist demise of authorial intention through an automated author that generates a series of words that appear to make sense to a reader.[27] But one of the most common problems with ChatGPT is that this appearance of thinking is almost entirely formal—the words seem like they make sense, but an attentive reading of ChatGPT's output reveals a complete lack of fidelity to "truth" or coherence. The words produced by AI do not "mean" anything in and of themselves, and any meaning in them comes more from a projection by the reader than anything in the text—an argument advanced by Bender and her colleagues, if one that is heavily debated.[28] The lack of intended meaning

23. See, for example, Steven Levy, "AI Is a Black Box. Anthropic Figured Out a Way to Look Inside," *Wired*, May 21, 2024, www.wired.com/story/anthropic-black-box-ai-research-neurons-features; Devin Coldewey, "WTF is AI?" *TechCrunch*, June 1, 2024, https://techcrunch.com/2024/06/01/what-is-ai-how-does-ai-work. Compare this to the claims of Offert and Dhaliwal, "The Method of Critical AI Studies."

24. See Adrian Mackenzie, *Machine Learners: Archaeology of a Data Practice* (MIT Press, 2017).

25. See Anthony McCosker and Rowan Wilken, *Automating Vision: The Social Impact of the New Camera Consciousness* (Routledge, 2020) and James E. Dobson, *The Birth of Computer Vision* (University of Minnesota Press, 2023).

26. Emily M. Bender, "Opening Remarks on 'AI in the Workplace: New Crisis or Longstanding Challenge,'" *Medium*, October 1, 2023, https://medium.com/@emilymenonbender/opening-remarks-on-ai-in-the-workplace-new-crisis-or-longstanding-challenge-eb81d1bee9f.

27. Compare this to N. Katherine Hayles, "Inside the Mind of an AI: Materiality and the Crisis of Representation," *New Literary History* 53, no. 4 (2022)/54, no. 1 (2023): 635–66, https://muse.jhu.edu/article/898324.

28. Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell, "On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? 🦜," FAccT '21: Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency, Virtual Event, Canada, March 3–10, 2021 (ACM), doi.org/10.1145/3442188.3445922.

from ChatGPT has led evolutionary biologist Carl T. Bergstrom and computational biologist C. Brandon Ogbunu to claim that, drawing on philosopher Harry Frankfurt, the output of LLMs like ChatGPT—filled with falsehoods and mistakes dismissed with the euphemism "hallucination"—are best described as "bullshit."[29]

I feel that Bergstrom and Ogbunu are generally correct in their use of bullshit to describe AI, even though there are some limitations to this characterization—not the least is, following much of what I've been arguing so far, it requires the "truth" of any statement made by an AI to remain at the level of human judgment, intention, and interpretation. As Frankfurt describes it, bullshit "need not be false" and

> it differs from lies in its misrepresentational intent. The bullshitter may not deceive us, or even intend to do so, either about the facts or about what he takes the facts to be. What he does necessarily attempt to deceive us about is his enterprise. His only indispensably distinctive characteristic is that in a certain way he misrepresents what he is up to.[30]

For Frankfurt, bullshit defines statements that cannot be evaluated as true or false, in part because the speaker has no real interest in the truth or falsity of what they say or claim. Thus, Bergstrom and Ogbunu hesitate slightly with their use of bullshit to describe AI—for something to be bullshit there should be a bullshitter, and they are uncertain that they want to attribute this level of intentional agency to AI. Yet, AI certainly has a form of agency, and, more importantly, this agency, *pace* Frankfurt, "in a certain way . . . misrepresents what [the AI] is up to." At the same time, and as I'll turn to explore in more detail momentarily, truth and falsity, with AI, depends mostly on statistical measures of probability. What is "true" is what is statistically probable, not what corresponds to "truth" in the sense of rationalistic, symbolic argumentation. Bullshit, and its distinction from a lie, is about intention for Frankfurt—or, at least, a lack of intention relative to the value of truth. This presumes a statement produced by AI can be judged according to the rules of logic and reference that follow philosophers of argumentation from Ludwig Wittgenstein through Stephen Toulmin to Frankfurt. But, if what I've been arguing is the case—that generative AI provides an environment for human sensation that distinguishes it from the symbolic logic of GOFAI (which *could* be evaluated in terms set by philosophers of logic and language)—then calling a statement made by ChatGPT "bullshit" misses the point *because this claim depends on a normative understanding of truth and symbolic reference*. It is this violation of the rules of symbolic reference and how "meaning" and argument are constructed that, in part, leads to the perplexity generated by generative AI.

The rest of this essay now moves into a discussion of some of the statistical methods that may provide an alternative to arguments about truth, argumentation, and logic for

29. Carl T. Bergstrom and C. Brandon Ogbunu, "Chat GPT Isn't 'Hallucinating.' It's Bullshitting," *Undark*, April 6, 2023, undark.org/2023/04/06/chatgpt-isnt-hallucinating-its-bullshitting. See also Michael Townsen Hicks, James Humphries, and Joe Slater, "ChatGPT is bullshit," *Ethics and Information Technology* 26, no. 38 (2024), doi.org/10.1007/s10676-024-09775-5.

30. Harry G. Frankfurt, *On Bullshit* (Princeton University Press, 2005), 54.

understanding the "perplexity" of AI. I'm also interested in other forms of AI besides LLMs, even though LLMs will appear throughout the next section because of their dominance in the current wave of AI, but my claims also relate to AIs that generate images. Much of what I've reviewed above needs to account for the qualitative and epistemological differences between language and images, which is a theme that characterizes a significant line of thinking in art history, media studies, and aesthetics.[31] But they also need to account for the separation of these words and images from the statistical logic of an AI. Or, the visual dimensions of AI provide an additional problem to what I've reviewed so far—the input and output of AI systems rely on words, and AI images are inevitably translated into language or, at the very least, groupings based on symbolic differentiations (albeit symbolic differentiations that have little to do with symbolic logic) that are themselves *statistical*.

I hesitate to fully embrace terms like "language" because the appearance of language as an input and an output is a distortion. I even hesitate to use "symbolic," although my use of this term presumes that there can be symbols that do not correspond to human sensation and meaning-making.[32] Any emphasis on the visual and linguistic dimensions of AI *also* needs to contend with the fact that the capacity of a computer to manipulate symbols is itself a distortion—"eyewash," in the words of media theorist Friedrich Kittler—that obscures how all computation reduces to voltage differences, not symbols or images.[33] While I generally agree with Kittler on this point, I don't fully reject symbols or images, as long as we accept that what constitutes symbols or images for an AI do not correspond to what these may mean for a human being. As I'll detail, text-based language models like ChatGPT rely entirely on the transformation of words into computations, into numerical "tokens" that are evaluated according to particular statistical weights in relation to other "tokens." If what we see on the screen appears in the form of words or images, how do we even address the aesthetics of computation? If we define aesthetics here in terms of *aesthesis*, then we foreground the relation between computers and sensation, which is never something that can ever be truly apprehended even though generative algorithms are shaping and determining the capacities of human sensation.

Accounting for these numerical "tokens" as a replacement for words and images requires a detour through the concept of *perplexity* as it is used in information theory and AI—as a specific statistical measurement of probability and uncertainty. The aesthetics of perplexity, when it comes to AI, is less about the strangeness of the images we see on a screen, or the awkwardness of the words generated by ChatGPT, than it is about the ability to sense—or the lack of an ability to sense—how computers have produced a space for knowing what a human being is and what a mind thinking can

31. See for example W.J.T. Mitchell, *Iconology: Image, Text, Ideology* (University of Chicago Press, 1986).

32. Offert, "Can We Read Neural Networks?" See also discussions of Félix Guattari's "a-signifying semiotics." A good summary of Guattari on this topic can be found in Gary Genosko's *Critical Semiotics: Theory, from Information to Affect* (Bloomsbury Academic, 2016).

33. Friedrich A. Kittler, *Gramophone, Film, Typewriter*, trans. Geoffrey Winthrop-Young and Michael Wutz (Stanford University Press, 1999), 1; also see Friedrich A. Kittler, "There Is No Software," in *The Truth of the Technological World: Essays on the Genealogy of Presence*, trans. Erik Butler (Stanford University Press, 2014), 219–29.

do. I'm now going to address perplexity as a statistical measure, and then conclude by describing how this statistical logic relates to a distinction between images, words, and numerical "tokens."

EXPERIENCING THE STATISTICAL

As mentioned above, the truth value of a statement made by AI comes less from any correspondence with reality, any accuracy in describing or representing what is "true," and is instead a result of what is statistically probable. This means that an inquiry into the politics and aesthetics of AI must take a detour through the epistemology of statistics,[34] as it is only through examining the logic of statistical reasoning that we can grasp just what decisions are going on that lead to the words and images produced by AI. As Justin Joque has argued, statistics operate both in terms of a mathematical and metaphysical logic, legitimating many contemporary transformations in capitalism and appealing to data and probability as the "truth" of the contemporary world.[35] Yet, this is not inherently the case with AI, and is also one of the reasons I started the above section with a reference to GOFAI and symbolic logic. What differentiates today's AI from earlier models of intelligence is this deferral to statistical reasoning—GOFAI was less about probability than about the proper integration of symbolic logic into how a computer senses and responds to a range of input. But the statistical logic of generative AI isn't completely new; the kinds of statistical measures used today derive from the history of information theory and cybernetics. In fact, we might suggest that the problem space of generative AI is distinct from AI in the past but is contiguous with the space of telecommunications and cybernetic control—even though, as numerous writers have argued, this space of information and cybernetic control has directly shaped many of the theoretical paradigms that guide the humanities and social sciences in the wake of World War II.[36] One measurement of statistical probability that originated in information theory and is central for the evaluation of AI today is termed *perplexity*.

In information theory, perplexity is a measurement of uncertainty. In one attempt to explain what perplexity refers to, without recourse to statistical formulas, mathematician and Meta data scientist Aaron Schumacher explains that perplexity is "an interesting measure of how well you're predicting something." If you want to, say, predict a sequence of numbers, each of which is between one and six, and you were using a six-sided die, you'd have a one in six chance that the die would be correct each time. The perplexity of the die would be six. "The perplexity of whatever you're evaluating, on the data you're

34. I'm inspired here by the work of Brian Rotman and Ian Hacking, both of whom have worked to study the history and semiotics of mathematics and statistics.

35. Justin Joque, *Revolutionary Mathematics: Artificial Intelligence, Statistics and the Logic of Capitalism* (Verso, 2022).

36. See for instance Bernard Dionysius Geoghegan, *Code: From Information Theory to French Theory* (Duke University Press, 2023) and Poornima Paidipaty, "'Tortoises all the way down': Geertz, cybernetics and 'culture' at the end of the Cold War," *Anthropological Theory* 20, no. 1 (2020): 97–129.

evaluating it on," Schumacher explains, "sort of tells you 'this thing is right about as often as an x-sided die would be.'"[37]

Perplexity is directly related to how Claude Shannon theorized "entropy" in his "A Mathematical Theory of Communication."[38] As Warren Weaver wrote in his explanation of Shannon's essay, Shannon's entropy is "associated with the amount of freedom of choice we have in constructing messages," a measure that Weaver equates with "information" more generally.[39] For Shannon, this measurement was entirely technical, mathematical, and related to a problem he had worked on at Bell Telephone Labs: how to separate signal from noise in communication systems.[40] As Shannon memorably wrote in "Mathematical Theory," even if communication is intended to transmit a message with meaning from one to another, "These semantic aspects of communication are irrelevant to the engineering problem. The significant aspect is that the actual message is one *selected from a set* of possible meanings."[41] The mere existence of semantic meaning, for Shannon, presumes a prior technical fidelity of communication, a technical fidelity that requires a finite number of potential options to be communicated. Both entropy and perplexity are measurements of this set of possibilities and require a medium of communication—usually a language—that is limited and constrained.

Now, it's important to foreground the initial task of Shannon's information theory and how this relates to some of the problems that persist today: a quantitative measure for the reduction or management of *noise*. As Shannon notes at the beginning of his paper, "The fundamental problem of communication is that of reproducing at one point either exactly or approximately a message selected at another point."[42] But, given the material reality of transmission, almost any message is disrupted by noise, which includes, in Weaver's interpretation of Shannon, "distortions of sound (in telephony, for example), or static (in radio), or distortions in shape or shading of picture (television), or errors in transmission (telegraphy or facsimile), etc."[43] The Mathematical Theory was an attempt to statistically understand transmission to handle problems like the poor quality of a telephone call or static in a television image. What is the boundary between a message making sense—in which the technical quality of a message is good enough to be interpreted—and a message that dissolves into noise?

Shannon's theory is easily applied to tasks related to language. Within any language, specific sounds, phonemes, symbols, and letters have a limited relation to the sounds, phonemes, symbols, and letters that follow. In English, it is not the case that, say, all

37. Aaron Schumacher, "Perplexity: what it is, and what yours is," *plan ➔ space*, September 23, 2013, https://planspace.org/2013/09/23/perplexity-what-it-is-and-what-yours-is.

38. Claude E. Shannon, "The Mathematical Theory of Communication," in Claude E. Shannon and Warren Weaver, *The Mathematical Theory of Communication* (University of Illinois Press, 1949).

39. Warren Weaver, "Recent Contributions to the Mathematical Theory of Communication," in *The Mathematical Theory*, 13.

40. See David A. Mindell, *Between Human and Machine: Feedback, Control, and Computing before Cybernetics* (Johns Hopkins University Press, 2002), 319–20.

41. Shannon, "The Mathematical Theory," 31.

42. Shannon, "The Mathematical Theory," 31.

43. Weaver, "Recent Contributions," 8.

twenty-six letters of the alphabet are equally probable at all moments. A Q is almost always followed by a U, for instance. At a statistical level, the probability that some other letter than U follows a Q is exceptionally low. The same can be said for almost any other sequence of letters, sounds, and even words, even though calculating the probability of words is far more difficult than letters.[44] Shannon's interest was calculating, statistically, how to measure the probabilities of specific symbols, letters, and words used in communication, and how these can be transmitted given the engineering specificities used in a specific medium.

There are countless engineering problems that descend from Shannon, and perplexity is one measure that came from those who followed in his wake. Perplexity was initially proposed by information theorist Frederick Jelinek in the 1970s, when Jelinek was working for IBM on problems related to automated speech recognition by computers. Jelinek defined perplexity as a "measure of the difficulty of the recognition task."[45] The larger the perplexity, the less likely a model will be correct in its predictions, which, for Jelinek, meant the less likely a computer would be able to recognize the speech of an individual and convert that speech into text.

Today, with generative AI, perplexity has become a measure of how well a language model can answer questions or respond with language that makes sense to a human being. From within computer science and information theory, the centrality of perplexity in evaluating language models has been regularly criticized—as with Shannon's original Mathematical Theory, perplexity as a measure cannot account for semantic changes in the meaning of words, for instance. Perplexity also cannot theorize a particular problem in the development of AI language models: almost all generative AI language models get worse over time, a fact that measures of perplexity cannot account for, a problem that the general logic of scale and data also, bizarrely, cannot account for.[46] This may be a result of a problem that comes from arguments like Parisi's—even though generative AI deals with uncomputable numbers, statistical measures of probability and uncertainly like entropy and perplexity inherently wrangle uncomputable possibilities into a finite, computable problem space. As LLMs get larger, they simply cannot deal with managing both the finitude of "sense" and also uncomputability. For the most part, these are common objections, as LLMs often have a difficult time differentiating between homonyms and cannot easily account for polysemy—a problem that comes from the specific way that LLMs convert words into numbers that I'll detail below. Nonetheless, perplexity is a general, statistical measure used to evaluate the general quality of LLMs.

44. This is related to Weaver and Shannon's interest in "Basic English," a project intended to "establish English as the standard language for communication worldwide" that was funded by the Rockefeller Foundation. Basic English used a limited corpus of words and, as described by Bernard Geoghegan, the creators of Basic English were aided "by a statistical analysis of English" they used to devise "a combinatory matrix for assembling complex meanings out of elemental combinations." See Geoghegan's *Code*, 28.

45. Frederick Jelinek, "The Dawn of Statistical ASR and MT," *Computational Linguistics* 35, no. 4 (2009), 487n3.

46. See Iris Luden, Mario Giulianelli, and Raquel Fernández, "Beyond Perplexity: Examining Temporal Generalization in Large Language Models via Definition Generation," *Computational Linguistics in the Netherlands Journal* 13 (2024): 205–32.

One particularly useful example of how perplexity is employed in evaluation comes from a common dataset used to assess LLMs, named LAMBADA, which stands for LAnguage Modeling Broadened to Account for Discourse Aspects. LAMBADA shows how both the concerns of Shannon persist, but also how Shannon's bracketing of semantics is, today, a major issue when the goal is not only the fidelity of a message being transmitted, but the generation of text that is supposed to mean something to a human being. LAMBADA, in the words of its creators, "proposes a word prediction task where the target item is *difficult to guess* (for English speakers) when only the sentence in which it appears is available, but becomes *easy* when a broader context is presented."[47] LAMBADA provides a series of examples of word completion tasks that require semantic and contextual awareness, and uses these tasks to train and evaluate an AI model to predict words that would otherwise be difficult or, from a statistical perspective, highly improbable.

For instance, one of the LAMBADA prompts is the sentence "Do you honestly think that I would want you to have a _____?" If given only this sentence, it would be highly unlikely that anyone—human or LLM—could correctly complete this task. The sheer number of possibilities that could reasonably complete this sentence is massive. The correct word that completes this sentence is, however, *miscarriage*, which is "easy" to answer only if one is provided the broader context that precedes this sentence, which is a passage of dialogue between two individuals: "'Yes, I thought I was going to lose the baby.' 'I was scared too,' he stated, sincerity flooding his eyes. 'You were?' 'Yes, of course. Why do you even ask?' 'This baby wasn't exactly planned for.'"[48]

As the creators of LAMBADA note, this example is an exemplary one, in part because another word—*abortion*—could almost complete the prompt sentence with the same broader context. Yet, solving this problem "requires a combination of linguistic skills ranging from (morpho)phonology (the plausible target word *abortion* is ruled out by the indefinite determiner *a*) through morphosyntax (the slot should be filled by a common singular noun) to pragmatics (understanding what the male participant is inferring from the female participant's words), in addition to general reasoning skills."[49]

I find the specific semantic content of this example to be strange—it provides yet another example of benchmarks or thought experiments in the history of computing, like Turing's "Imitation Game" or Searle's "Chinese Room," that invoke a gendered or racialized example for why a computer can or cannot think.[50] Even though there are

47. Denis Paperno, et al., "The LAMBADA dataset: Word prediction requiring a broad discourse context," in *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), Berlin, Germany* (Association for Computational Linguistics, 2016), 1525, https://aclanthology.org/P16-1144.pdf.

48. Paperno et al., "The LAMBADA Dataset," 1527.

49. Paperno et al., "The LAMBADA Dataset," 1531.

50. Turing's "imitation game" and its relation to gender has been widely discussed, from Jack Halberstam's "Automating Gender: Postmodern Feminism in the Age of the Intelligent Machine," *Feminist Studies* 17, no. 3 (1991): 439–60, www.proquest.com/docview/233180777?sourcetype=Scholarly%20Journals to David Golumbia's "Computation, Gender, and Human Thinking," *differences: A Journal of Feminist Cultural Studies* 14, no. 2 (2003): 27–48. The function of racial difference in Searle's "Chinese Room" experiment has been less discussed but Lydia H. Liu's "After Turing: How Philosophy Migrated to the AI Lab," *Critical Inquiry* 50, no. 1 (2023): 2–30, provides

over ten thousand passages in the LAMBADA dataset, this example, which is discussed multiple times by LAMBADA's creators in the paper that introduces their dataset, ends up presuming "thought" to be linked with the ability to infer several highly gendered relations from a brief passage.

The problems don't end with this invoking of gender. LAMBADA's dataset is derived from "unpublished" novels—I put this term in quotes because it is the term the creators of LAMBADA use to describe its sources, though this characterization is incorrect. LAMBADA's creators argue that the

> fact that it contains unpublished novels minimizes the potential usefulness of general world knowledge and external resources for the task, in contrast to other kinds of texts like news data, Wikipedia text, or famous novels. The corpus, after duplicate removal and filtering out of potentially offensive material with a stop word list, contains 5,325 novels and 465 million words.[51]

The specific dataset LAMBADA uses is the "BookCorpus," which is comprised not of "unpublished" works, but is a dataset of seven thousand self-published novels taken illegally from the ebook distribution platform Smashwords. The BookCorpus was the original dataset used by OpenAI to train its initial GPT model. In other words, LAMBADA is, like so many things related to the training of LLMs, dependent on the theft of a massive amount of intellectual property, and also strangely privileges "complete" novels that are not part of mainstream cultural production and distribution, approaching the countless platforms that exist for user-generated content as if this material is without social and economic value, even presuming these novels exist without readers or a public (else they would be "useful").[52]

Regardless, without fully bracketing these criticisms, the goal with LAMBADA is to generate a corpus of training data with a relatively high perplexity. While, to use the examples provided in the original LAMBADA paper, a standard LLM may have a perplexity of anywhere between 149 and 566, the LAMBADA set of tasks has instead a perplexity of between 769 and 16318. The point here is to create language models that do not rely only on a preceding word to make a prediction, but a much broader context. If an LLM can "remember" and make reasonable inferences based on context, then perplexity would be reduced even when using a large corpus of words.

So far, however, this explanation of LAMBADA and perplexity may not seem to follow what I've argued in the first part of this essay. The goal would seem to be to account for semantics, for meaning, for human perceptions about language and symbolic

a critique of race in Searle's work, and the artist Candice Lin's 2018 work *On Being Human (The slow erosion of a hard white body)* is built around a critique of Searle's presumptions surrounding race, otherness, and norms of "transparent" communication built into the Chinese Room thought experiment.

51. Paperno et al., "The LAMBADA Dataset," 1526–28.

52. Richard Lea, "Google swallows 10,000 novels to improve AI's conversation," *The Guardian*, September 28, 2016, www.theguardian.com/books/2016/sep/28/google-swallows-11000-novels-to-improve-ais-conversation. The original paper defining the BookCorpus dataset is Yukun Zhu, et al., "Aligning Books and Movies: Towards Story-Like Visual Explanations by Watching Movies and Reading Books," IEEE International Conference on Computer Vision (ICCV), Santiago, Chile, 2015.

reference. And, in fact, we could argue that perplexity, to draw on some of the jargon of semiology, is about narrowing the possibilities for a paradigmatic axis of a sentence, the signifier's capacity to be substituted with other signifiers, to ensure the coherence of the syntagmatic axis. But the difference here emerges from how the "memory" of an LLM works and how it makes inferences from a broader context.

In LLMs that are called "transformers," like ChatGPT, "remembering" happens through a mechanism referred to as "attention" in machine learning.[53] Intended to replicate, in a superficial way, how human attention works, what an LLM does with "attention" transforms words into specific numerical representations, or "tokens," and then gives numerical weights to these tokens based on a "word embedding table," which provides the probability of any co-occurrence of tokens. The transformer model was proposed by engineers at Google, and was used to create translation software, using datasets of 4.5 million sentences paired in English and German, and then a larger dataset, comprised of 36 million sentence pairs in English and French. For English and German, words were converted into a "vocabulary" in which both sources and targets contained approximately 25,000 tokens. My supplying of these details is intended to highlight the following: generative AI systems are, in the end, about transforming something into a numerical representation, a "token," and then correlating one number with another number probabilistically. Even though perplexity—and many of the tasks that descend from Shannon—would seem to be about letters, phonemes, words, and so on, in the end everything must necessarily be converted into a numerical token. All forms of representation must inherently reduce to these tokens, and the relations between tokens do not come from symbolic reference, but from how the LLM represents these relations to itself.[54]

While I've been focusing on LLMs, as these statistical measures originate with LLMs, the ability to reduce any and all input to a token is also essential for AI image generators. Almost all AI image generators—at least the most popular ones today, including OpenAI's DALL-E and Stability AI's Stable Diffusion—depend on what's termed a "diffusion model." Diffusion models are, at a basic level, trained to remove noise from visual images, and exceptionally basic predecessors to diffusion models can be found in software like Adobe Photoshop with "reduce noise" filters. To train a diffusion model, one begins with an image and then adds Gaussian noise to it, and then the diffusion model works to automatically reverse this process. The AI would be able to "repair" images through the ability to predict gaps, errors, and distortions in the image itself—a problem not overtly distinct from the problems addressed by Shannon. Yet, a diffusion model is not only about removing noise, but, again, about predicting probability distributions between tokens. In the paper that provides the mechanism that OpenAI used to develop DALL-E, the authors note that correlating individual pixels to tokens would "require

53. Ashish Vaswani, et al., "Attention is all you need," *NIPS '17: Proceedings of the 31st International Conference on Neural Information Processing Systems* (Curran Associates Inc., 2017), 1–11.

54. See Fabian Offert and Peter Bell, "Perceptual bias and technical metapictures: critical machine vision as a humanities challenge," *AI & Society* 36, no. 8 (2021): 1133–44.

an inordinate amount of memory for high-resolution images."[55] Instead, DALL-E was initially trained on a method to compress 256x256 RGB images into 32x32 grids of image tokens, each of which could be one of 8192 values. These image tokens were correlated with captions that could be a maximum of 256 tokens long with a vocabulary of 16,384. While we see a link between text and image, what DALL-E is actually doing is converting text into token and image into token—which, importantly here, ends in a "single stream of data."[56]

The point I'm making with all of this—moving from perplexity to some of the details of how tokens work in LLMs and diffusion models—is to highlight how, in any generative AI system, what we're actually experiencing (reading words on a page, inputting text into a prompt, looking at an image) has almost nothing to do with what the computer is actually doing. Instead, from the computer's perspective, letters, words, and pictures, the parts of a sentence and the formal elements of an image, all reduce to a single object—a "token"—that correlates to a number and its probability in relation to other numbers.

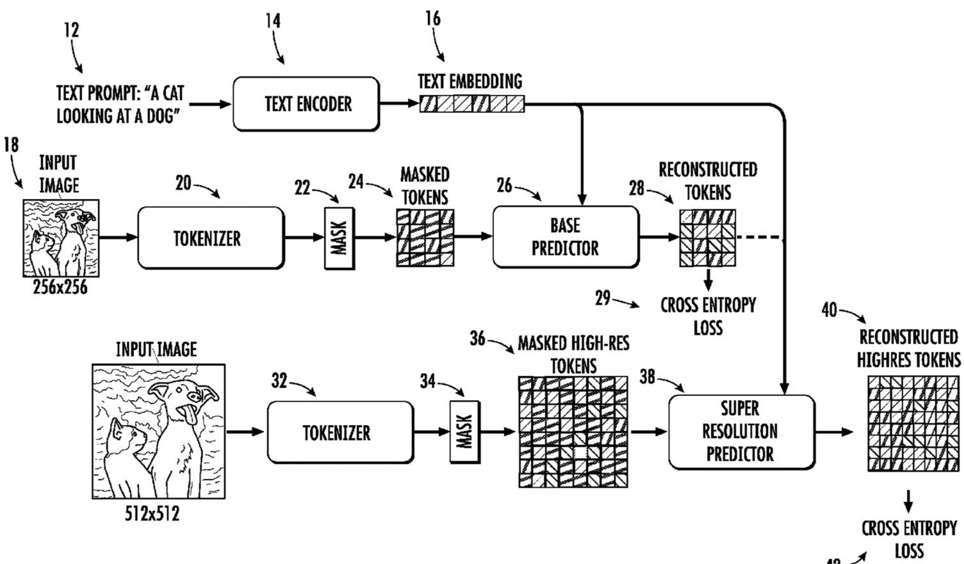

Diagram from Google's patent for "Text-to-image Generation via Masked Generative Transformers" (WO 2024/130137 A1). What Google proposes in this patent is not exactly the same as the underlying workflow for DALL-E 2, for instance, but is representative of what I've been discussing here. In this diagram, we can see how images are converted into "tokens" that have little to do with what humans see, correlated with text prompts that are also "encoded" as tokens, used to eventually "reconstruct" an image out of correlated tokens. Here, Google is specifically accounting for problems of resolution given how image generators are trained with images that are of much lower resolution than what human users would likely want from AI.

55. Aditya Ramesh, et al., "Zero-Shot Text-to-Image Generation," *ICML '20: Proceedings of the 37th International Conference on Machine Learning, Online* (MLR Press, 2020), 2.

56. Ramesh, et al., "Zero-Shot Text-to-Image Generation," 4.

This observation, of course, is not particularly unique—variations of this divide between human experience and the technical operation of a computer can be found throughout the theorization of digital culture.[57] But much of this divide, in the past, had to do with distinctions between source code and object code, pixels on a screen and the voltages of a processor. Here, however, what seems to me to be essential is the conversion of all cultural forms into tokens that can be arranged according to probability. Or, the conversion of all values and expression into a single table of mathematical correlations. We may see—in the output of any of these systems—what appear to be words and images. Benchmarking techniques like LAMBADA may seem to refer to typical semiological problems of meaning, reference, syntagm and paradigm. But this depends on the reduction of all sensory data—here, text and image, but obviously we're not limited to only correlations between text and image—to numerical tokens that can be processed as a single stream of data.

CONCLUSION: IMAGE, TEXT, NUMBER

For W. J. T. Mitchell, a classic problem of Enlightenment thought was that language and imagery ought to be "perfect, transparent media through which reality may be represented to the understanding." But, in the centuries since this Enlightenment ideal, critical thought has instead imagined language and images as "enigmas, problems to be explained, prison-houses which lock the understanding away from the world." In making this move, images gradually became a "kind of language," and differences between image and text deferred to ideological oppositions, actual pictures disappearing from contemplation.[58] My argument in this essay would probably seem, for someone like Mitchell, to be following closely in the steps of a range of previous claims about the visual dimensions of digital media. "If the Desert of the Real is," Mitchell claims, "just numbers, then we can take some comfort in the fact that Plato already made this point over two thousand years ago, and he still thought the only moral to be drawn from it was that we had to go on living in this world of shadows, illusions, and images . . . "[59]

In a certain respect, this is true—and follows countless other differentiations between experience and reality, phenomena and noumena, that can be woven through the entire history of Western thought. But I think there's a particularly strange logic going on that Mitchell's dismissal does not take into account. For the history of critical thought, going back to Plato, the surface phenomena is where we can locate ideology, while the underlying "reality" *is* reality. What we see here is not "just numbers," but numbers that are statistically correlated tokens. And what I want to suggest is that *it is in the mathematical operations beneath appearances where we can locate ideology*—and it

57. For instance, Kittler, "There is No Software;" Wendy Hui Kyong Chun, "On Software, or the Persistence of Visual Knowledge," *Grey Room* 18 (2005): 26–51; Alexander R. Galloway, "Language Wants To Be Overlooked: On Software and Ideology," *Journal of Visual Culture* 5, no. 3 (2006): 315–31.

58. Mitchell, *Iconology*, 8.

59. W.J.T. Mitchell, *Image Science: Iconology, Visual Culture, and Media Aesthetics* (University of Chicago Press, 2015), 59.

is this argument that we can find in the best critiques of statistics and mathematics today, in the work of people like Brian Rotman, Joque, and Parisi.[60] It is not that mathematics is "reality," be it the reality of computation or reality *tout court*. But that, in the reduction of human experience to number, the ideological functioning of capitalism and digital culture are transformed into an "objective" quantity that then, through the apprehension of massive quantities of data that are beyond the human ability to sense or understand, shape the possibilities for sensation, experience, and agency, without ever entering into a human ability to understand or even intervene. What's at stake is *not* that numbers are "reality," but that today's ideologies, encoded in the operations of numerical tokens, *are what remain invisible.* It is almost impossible to critique this ideology because these ideologies remain almost entirely veiled to human experience and understanding.

For Mitchell, "ideology" refers to that which produces "the difference between the 'natural' and 'conventional' sign; the distinction between an illiterate viewer, who can see what images represent, and a literate reader, who can see through the image to something else (typically, a text)."[61] What I hope I've shown above is how, with AI, it's not just that one inputs text and gets a new text in return, or an image derived from text, but that to make these translations and transformations words and images must be rendered equivalent through their technical existence as "tokens" that are correlated on a table of all possible values. Perplexity is a name for the statistical limitation of these possible values. The more possibilities, the greater the perplexity. The goal of "better" AI is to reduce these possibilities, to reduce the perplexity. It is not that, with the digital, the "reality" of the world is numerical, but that a filtering of all other values through number and probability is the ideological mechanism upon which all other sensation and experience relies.

I think that the problems I surveyed in the first part of this paper are a result of the technical issues discussed in the second. The goal of perplexity as a statistical measure is the reduction of perplexity as an experience, the creation of a world that is predictable, controlled, and without the messiness and confusion that comes from the unexpected, the improbable, the human. I'm not suggesting—completely—that digital media and AI don't need people, but that the transformation of ideology into a statistical operation that itself remains obscure makes it feel as if digital media don't need people, a feeling that would resonate with many of the broad forms of disempowerment woven throughout society and culture today. If an aesthetics of perplexity, in the sense that this special issue is mostly addressing it, is about a mismatch or a conflict between observer and observed, an opening of the possibilities of sense that does not resolve, then the statistics of perplexity are there to reduce the possibilities of this experience. I think that, in many ways, we know this—but are at a loss of how to deal with it because the ideological

60. Along with Joque, *Revolutionary Mathematics*, and Parisi, *Contagious Architecture*, see Brian Rotman, *Mathematics as Sign: Writing, Imagining, Counting* (Stanford University Press, 2000), which explicitly emphasizes the figurative and semiotic dimensions of mathematics, not the "natural" existence of mathematical signs.

61. Mitchell, *Image Science*, 43.

operations happening with computation are beyond human experience. An environment brought about through AI is a world, ideally, without confrontation, as minor and mild as these confrontations may be. ■

GRANT BOLLMER is an associate research professor in the department of communication at the University of Maryland, College Park. His most recent books include *The Affect Lab: The History and Limits of Measuring Emotion* (University of Minnesota Press, 2023) and, with Katherine Guinness, *The Influencer Factory: A Marxist Theory of Corporate Personhood on YouTube* (Stanford University Press, 2024).